

DAiSEE: Towards User Engagement Recognition in the Wild

Abhay Gupta
Microsoft

abhgup@microsoft.com

Vineeth N Balasubramanian
Indian Institute of Technology Hyderabad

vineethnb@iith.ac.in

Abstract

Recognizing user engagement plays an important role in several contemporary vision applications including advertising, healthcare, and e-learning. However, the lack of any publicly available dataset to address this problem severely limits the development of methodologies that can help make progress to address this challenge. In this paper, we introduce DAiSEE, a free-to-use large dataset comprising of 8000 video snippets captured from 95 users for recognizing user engagement in the wild. Baseline results using standard feature extraction and classification methods show the difficulty of working with this dataset. We believe that DAiSEE would provide the research community with challenges in feature extraction, context-based inference as well as the development of suitable machine learning methods that can consider annotator statistics in the training process itself, thus providing a springboard for further research.

1. Introduction

The progress from lab research to consumer technologies in classical problems of recognition in computer vision has largely been possible over the last decade due to the availability of large-scale datasets which are made available to researchers and industry practitioners alike. The ImageNet [9] and PASCAL VOC [13] challenges spearheaded the object recognition revolution, while the recent availability of datasets such as Microsoft COCO [23] is fueling the development of newer methods such as in vision-to-language and language-to-vision efforts. Needless to say, the availability of more datasets for various problems and subproblems in computer vision will allow a greater impact in translating research methodologies to businesses and, eventually, a positive influence on users' lives. This work is an effort in this direction - to provide a dataset (and benchmark results) for a vision problem of contemporary relevance: user engagement recognition.

The recognition of user engagement is increasingly relevant to a digital world that floods users with various kinds of content, and it is useful for a system to be "aware" of

the user's engagement while providing content. For example, the recognition of user engagement is critical to various domains including advertising - web or television (is the viewer really watching this advertisement on television or on YouTube? is (s)he bored?); healthcare - as pertains to detection and interventions for children with learning or cognitive disabilities (does a child's engagement levels indicate a tendency for autism or ADHD? are there particular objects or events that interest or frustrate the child?); and or e-learning (which parts of a lecture are confusing for most students who watch it? how engaged are students in a video?). While there have been recent research efforts to develop methods for recognition of user engagement [37] [18], their datasets are not publicly available. Further, existing commercially available affective recognition systems that attempt to track user engagement work within constrained settings, and have limited use in real-world environments (illustrated further in Section 2). This work introduces DAiSEE, a dataset that aims to facilitate research and development towards user engagement recognition in the wild.

Understanding affective (emotional) states of a user, an important subarea of computer vision, has for a long time focused on datasets pertaining to the seven basic expressions: neutral, happiness, sadness, anger, disgust, surprise and contempt [24]. While recent efforts in the last few years have expanded datasets in this domain to cover emotions in terms of a dimensional representation [33] [28], the vast subtleties in emotions and affective states necessitate the development of datasets for specific objectives. Recent trends, including [33] [28], corroborate this approach to help progress towards tangible outcomes.

E-learning environments provide one of the best use cases for studying user engagement. With the accelerated growth of Massive Open Online Courses (MOOCs), there is a need to design more intelligent interfaces to simulate the interactions that occur between a teacher and students in a class. The main drawback of existing e-learning environments is that they do not provide real-time interactive feedback to students (or instructors, for that matter) during the content delivery process, as compared to traditional



Figure 1. Examples of images from DAiSEE. The dataset captures real-world challenges of recognizing user engagement in natural settings.

classroom learning. Surprisingly, MOOCs have a dropout rate of 91-93% [21], with the completion rate for the first assignment being around 45%. An online survey [7] lists the top ten reasons for dropouts from such platforms; poor course design was one of the reasons, which included components such as lack of proper student feedback, “lecture fatigue” in courses that had only video lectures, lack of proper course introductions and student frustration. Such reasons motivate the need to improve feedback mechanisms and make such platforms more interactive. Understanding user engagement at various junctures of the e-learning experience can help design more intuitive interfaces that further knowledge absorption by students and help decrease dropout rates, as well as personalize the learning experience. This paper seeks to address the aforementioned issues, by making available a dataset that captures user engagement during e-learning sessions in the wild (which can also be relevant to other applications such as advertising and healthcare, that have similar use case settings).

In this work, we introduce DAiSEE (**D**ataset for **A**ffective **S**tates in **E**-learning **E**nvironments) (Figure 1 shows sample images), which will be made publicly available to the vision community for further research. In particular, we focus on engagement, frustration, confusion and

boredom as the affective states considered in this work, all of which are relevant to user engagement and concomitant applications. Considering that such affective states are subtle, the annotations for this dataset was crowdsourced and provided along with this dataset. We also benchmark the performance of standard feature extractors and classifiers on this dataset to provide a baseline for further research.

The rest of the paper is organized as follows: We discuss the background and related work in Section 2. In Section 3, we introduce the dataset and its salient features. In Section 4, we benchmark a baseline performance on this dataset using recent improvements in deep learning techniques. Lastly, we summarize our analysis and suggest future directions with our dataset.

2. Background and Related Work

Determining the affective state of a user using computer vision and machine learning methods has been studied for over two decades now (please see [39] [31] for surveys). As mentioned earlier, until recently, most efforts focused on the seven essential expressions (neutral, happiness, sadness, anger, disgust, fear, surprise) and the facial action units connected with them, as in [38] [32]. Despite recent efforts to go beyond these basic expressions as well as model af-




Figure 2. Results of Affdex on videos from our dataset. The first entry on the text measures engagement. In image (a), the software detects that the subject is attentive (engaged), even though the subject is not. In subfigures (b) and (d), the software detects a highly engaged person as having minimal engagement (attentive). In subfigure (c), the software detects zero attention (engagement), even though the subject is engaged.

ffective states in terms of dimensions such as valence and arousal [14] [17] [33] [28], very limited work has been carried out in perceiving affective states related to user engagement, which has direct relevance on various applications as described before.

Hernandez et al. [18] were the first to attempt the problem of recognized user engagement. They modeled the problem of determining engagement of a TV viewer as a binary classification task, using multiple geometric features extracted from the face, and SVMs for classification. Considering the lack of a publicly available dataset, a custom dataset - very small and labeled by a single coder - was used in their work (which, unfortunately, is not publicly available too). Similarly, Whitehill et al. [37] attempted to automatically understand engagement in learning environments. They developed a custom dataset labeled by a few coders, which again is not publicly available. They experimented with feature extraction methods and classifiers and concluded that Support Vector Machines with Gabor features gave the best accuracy on recognizing user engagement on a scale of 4 levels (highest accuracy being $\approx 69\%$). Besides, the dataset was captured under constrained settings and did not capture the "in-the-wild" needs of real-world user engagement recognition.

The relevance of affective state recognition in real-world applications can be gauged by the rising number of commercial applications that attempt to address this challenge. Applications such as Emotient [11], Emovu [12], and Sightcorp [30] provide an estimation of comparable affective states (called attentiveness, for instance, in SightCorp) in their frameworks. On one hand, all these applications are constrained only to attentiveness/engagement and do not consider other related affective states, such as boredom or confusion (which we seek to capture in this work). More importantly, our studies with these applications showed that their performance on real-world videos is far below satisfactory, thus highlighting the need for a dataset that captures real-world conditions for further research. Figure 2 shows an example of the performance of Affdex [1] on videos from

our dataset and we see that the software shows a user to be attentive even if the user's eyes are closed or the user is looking away from the screen. We note from the image that the software tracks facial key points and correlates them with emotional and cognitive states. Other applications such as SightCorp [30] use the eye gaze of the subject as the sole determinant of the engagement level. While these are good starting points, the availability of a larger dataset is likely to promote better methods towards the end of reliable engagement recognition.

Why e-Learning? Subsequent to the exponential growth of MOOCs over the last few years, e-learning has received significant attention from several research groups. While machine learning methods have been used to personalize educational modules, diversify assessment methods and make personalized recommendations based on learner preferences and browsing patterns [6] [19] [3] [26], limited efforts have been attempted in understanding user engagement in the e-learning environment. More importantly, e-learning provides a well-defined context for developing this dataset, as well as studying the impact of effective recognition methods, thereby motivating us to choose this context for this dataset, without any loss of generality.

3. The DAiSEE Dataset

We now present the DAiSEE dataset that contains video recordings of 95 subjects in an e-learning environment, annotated with labels for engagement, frustration, confusion and boredom levels using crowdsourcing. The dataset captures "in the wild" settings typically seen in the real-world, and will be made publicly available, along with the individual annotations of the crowd, to facilitate open research. We now discuss the data collection procedure, the data annotation process followed, and the vote aggregation strategy used to create the dataset.

3.1. Data Collection

We use a full HD web camera (1920x1080, 30 fps, focal length 3.6mm, 78° field of view) mounted on a computer focusing on student users watching videos on the computer monitor. To simulate the e-learning environment, a custom application was created that presented a subject with 2 different videos (20 minutes total in length), one educational and one recreational to capture both serious and relaxed settings, which allow natural variations in user's engagement levels. To model unconstrained settings, the subjects had the option to scroll through the videos.

There are 95 subjects in the dataset belonging to the age group of 18-30, all of whom are currently enrolled students. The race of the subjects is Asian, with a male-to-female ratio of 37: 3¹. In total, 30 hours of recordings were captured. The resulting dataset has 8000 video snippets across 5 different locations and 2 different illumination settings. Each video snippet is 10 seconds long and encoded at 30fps.

The Hawthorne effect [16], also referred to as the observer effect, is a type of reactivity in which individuals modify or improve an aspect of their behavior in response to their awareness of being observed. This is a critical aspect of such a data capture setting and it is highly probable that the subjects may adapt their behavior to suit the objectives of the experiment. To diminish the effects of such a circumstance, the subjects were recorded without their knowledge. This helped in limiting the Hawthorne effect. To account for the privacy interests of every subject, at the end of the data capture, they were informed of the recordings and their consent obtained to carry out research work. In the event that a subject declined consent, the captured videos were deleted. Further, the anonymity of every subject is ensured by giving him/her a uniquely generated 3-digit id whose correspondence with the identity is not recorded anywhere.

3.2. Data Annotation

3.2.1 Class Labels

Motivated by recent work in intelligent tutoring systems [27], our dataset consists of labelings of four affective states related to user engagement, viz., engagement, frustration, confusion, and boredom. Recent work [5] has shown that the six essential expressions: anger, disgust, fear, joy, sadness and surprise [10] are not reliable in prolonged learning situations, since these basic expressions are prone to rapid changes. Each of the affective states is defined in four levels: (1) very low (2) low (3) high and (4) very high (similar to [37]). We followed this class label strategy, in particular, to avoid the "neutral" state since early experiments showed that crowd annotators often preferred

to choose "neutral" as a state when not sure. The aforementioned definition of the possible levels ensured that the annotators "took a stand" on the engagement level, which is essential for a robust dataset. We also note that as in [37], the above 4-level annotation can be trivially changed to 2-levels (*high* and *low*) as may be required for a given application setting.

3.2.2 Annotation Process

As mentioned earlier, DAiSEE is created by tapping into the potential of crowd annotators. Subtle affective states such as user engagement can be hard to label and can vary based on a viewer's discretion. Hence, we rely on the "wisdom-of-the-crowd"'s decision for our annotations in this dataset. The large amounts of data and easy availability of annotators on crowdsourcing platforms has resulted in many other newer computer vision datasets such as [9, 23, 34] to tap into the crowd for annotations. Although the annotators can be non-experts, it has been shown that repeated labeling of examples by multiple annotators produces high-quality labels [20,29,36]. In this work, we used CrowdFlower, a popular site for crowdsourced tasks, for the annotations, similar to [4]. CrowdFlower provides advanced quality control mechanisms, worker targeting and detailed reports on the final annotation results obtained. Other features of CrowdFlower that we used in this work are mitigation of bot labeling, priming of annotator to the specific task using reasoned test questions, and flagging labels of underperforming annotators.

For DAiSEE, we have four different affective states, each having four levels, as described earlier in this section. To obtain the votes, each annotator is presented with a video snippet and asked to vote. Annotators are presented with instructions on how to perform the task and illustrative examples to facilitate the process. Additionally, each annotator answers a standardized test question, that helps us remove the votes of underperforming annotators. For each video snippet, we get votes from 10 different annotators (which is comparable to other standard crowdsourced datasets such as [41] [15]).

3.3. Vote Aggregation

After obtaining the votes for all video snippets, we use the Dawid-Skene [8] strategy, which is widely used as a gold standard for vote aggregation methods, to aggregate annotator labels to obtain the single ground truth label for each video. Dawid-Skene is an unsupervised inference algorithm that gives the Maximum Likelihood Estimate of observer error rates using the EM algorithm. The algorithm has the following main steps:

1. Using the labels given by multiple annotators, estimate the most likely "correct" label for each video snippet.

¹More female subjects are planned to be captured in future, and will be appended to the dataset




Figure 3. Label distribution after applying the David-Skene aggregation method for all affective states across all 8000 videos.

2. Based on the estimated correct answer for each object, compute the error rates for each annotator.
3. Taking into consideration the error rates for each annotator, recompute the most likely “correct” label for each object.
4. Repeat steps 2 and 3, until one of the termination criteria is met (error rates are below a pre-specified threshold or a pre-specified number of iterations are completed).

The final ground truth label distributions after applying the aggregation strategy for the four affective states is seen in Figure 3. We now discuss the benchmark results obtained using standard methods on this dataset.

4. Benchmark Results

The videos in DAiSEE are captured with the frame of capture extending up to the bust of the user. This allows interested researchers to study the relevance of non-facial cues such as upper body postures on user engagement (for example, a laidback posture could indicate higher levels of boredom, while an upright posture could indicate high levels of engagement). However, in this work, we focus on obtaining benchmark results using the face extracted from each video frame in the dataset. We extract features from the cropped faces and use standard classifiers to obtain our baseline results.

4.1. Feature Extraction

We first perform face detection and alignment using the publicly available OpenFace [2] (which internally uses DLIB² tool). Once the faces have been aligned, we crop the faces in the frames of the videos. We then extract the following features from the face region of each video snippet:

²DLIB:<http://dlib.net>



Figure 5. Architecture of VGG Face

3D HOG Descriptors: We extract dense 3D HOG [22] from the cropped video and then process the features for a bag-of-words type representation. We apply k -means clustering on the descriptors with $k = 256$ clusters [35]. This results in a 256-dimensional frequency histogram of facial features. We then normalize the histogram of the features.

3D LBP Descriptors: We extract the LBP-TOP [40] features for every video. The features for all the orientations are concatenated for the feature representation of a video.

Deep Face Descriptors: We use VGG-Face [25], a Convolutional Neural Network-based feature extraction method, for these descriptors. The architecture of VGG Face is shown in Figure 5. We extract the features from the *fc7* layer to get the most generic features as the *fc8* representations are tailored for the dataset proposed in [25]. The extracted features are processed for a bag-of-words representation. We apply k -means clustering on the descriptors with $k = 50$ clusters. This results in a 50-dimensional frequency histogram. We then normalize the histogram of the features.

4.2. Classifiers

After all the features have been obtained, we use three classifiers, namely k -NN, SVM and Random Forests for classification. For SVM, we ran preliminary studies with four different kernels: RBF, Linear, Poly-3 and Poly-4 and found that RBF kernels performed the best among them. For k -NN we varied k from 1 to 100 and found that at $k = 49$, we obtained the best mean performance for the different feature extraction strategies. For random forests, we used a forest size of 150 with a minimum split of 2 where we empirically observed the highest average accuracy.

4.3. Performance Metrics

5-fold cross-validation is used to measure the generalization capabilities of the baseline methods. To ensure a fair comparison in case of any imbalance in the labels, we use the average classwise accuracy as the performance metric in this work (accuracy is measured for each class individually, and their average is presented), where the average is computed across the folds of cross-validation.

4.4. Results, Analysis and Discussion

Figure 6 presents the baseline results from our studies. Some example results are shown in Figure 7. We analyze

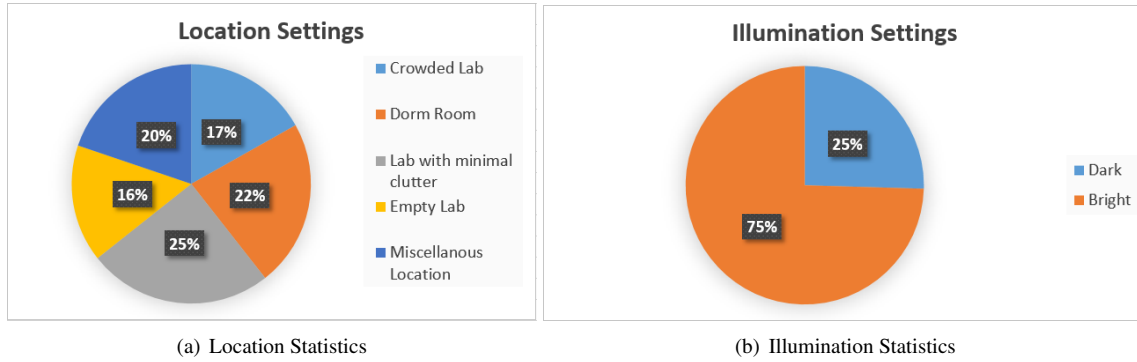


Figure 4. *Left*: Distribution of videos across 5 location settings *Right*: Distribution across illumination settings: low and high.

the performance of our benchmark results on DaiSEE, including the dataset’s behavior to changes in illumination and location settings, below.

We observe from Figure 6 that k -NN and random forests consistently outperform SVMs in our studies. In terms of feature extraction methods, 3D LBP performs better than 3D HOG and VGG-Face, although marginally. We believe that this is because we use a bag-of-words representation for 3D HOG and VGG-Face, which may not necessarily capture the characteristics of the video. Learning features in an unsupervised manner using deep learning architectures that are tailored to this specific problem is an important direction of work that could help improve the performance, considering the recent successes of deep learning.

In summary, we observe that the classwise accuracy performance is only marginally better than random (considering there are 4 classes), thereby showing the difficulty of working with this dataset. We note that the dataset captures the nature of real-world e-learning environments in an organic manner, with varying user poses, positions and background noise typically encountered in such settings. Methods that use geometric features (such as facial fiducials), facial action units, body pose and eye gaze may need to be explored to improve the performance w.r.t. the baselines shared in this work. It is also possible that our ground truth labels are noisy due to their crowdsourced nature. Considering annotator statistics during vote aggregation can help in improving ground truth labeling itself too.

4.5. Confusion Matrix

In order to better understand the results, we present the confusion matrix for one of the result configurations. Table 1 provides us with the confusion matrix for engagement with 3DLBP being the feature extraction method and random forests as the classifier. This setting was chosen because it provided the best results for engagement. This shows that the *Very High* state (label 4) provides the best results for classwise accuracy, showing that it is easiest to

Confusion Matrix		Obtained Labels			
		1	2	3	4
Actual Labels	1	417	453	508	714
	2	205	688	397	588
	3	123	452	913	756
	4	104	350	435	896

Table 1. Confusion Matrix with the configuration, Affective State: Engagement, Feature: 3D-LBP, Classifier: Random Forests

notice very high engagement.

4.6. A 2-class approach to understanding engagement

We studied the dataset from a modified performance metric based on a 2-class approach to understanding engagement. If the predicted label is 2, but the true label is 1, then we consider this instance to be half-correct, while computing accuracy. However, if the label is 3 or 4, we do not consider this instance to be correctly classified. Similarly, if the predicted label is 4, but the true label is 3, we consider this instance to be half-correct while computing accuracy. This approach considers the subtlety of the labels of the dataset as it is often very difficult to distinguish between high and very high engagement and vice versa, low and very low engagement. The results from this approach are shown in Figure 8. Once again, we see the superior performance of 3D-LBP, which is considered a good feature extraction methodology for faces.

4.7. Applicability to Real-world User Engagement Recognition

To test the usefulness of DAiSEE to real-world user engagement recognition, we ran models trained on the dataset on a new data stream from a web camera. Figure 9 shows the results of this study. In this experiment, an Intel Xeon E5 with two 8-core 1.2GHz CPUs with 64GB RAM is used and we are able to process up to 100 frames/min in real-

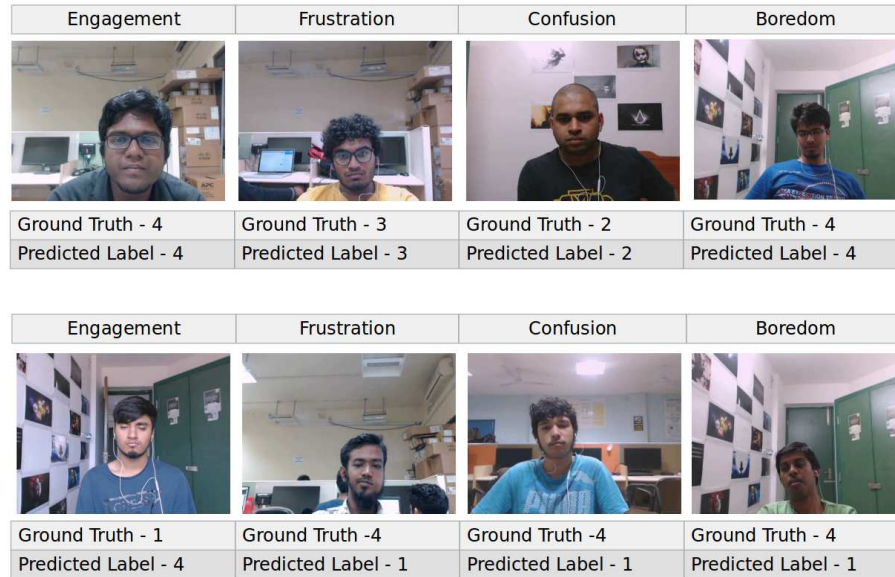


Figure 7. Example predictions from DAiSEE. All the results are generated using 3D-LBP using random forests classifier. The top image shows correct classifications and the bottom image shows misclassification results.

time video streams. The feature extraction and model validation are performed using a single core while the face detection and alignment use both the cores. While we do not have ground truth for these frames, the results show promise from a subjective evaluation.

5. Conclusions and Future Work

This paper introduces DAiSEE, a crowdsourced dataset that aims to provide the community with a dataset to capture user engagement in the wild. The proposed dataset has rich information including 4 different affective states, namely boredom, confusion, engagement and frustration across various location and illumination conditions and videos presented during the capture.

Our future work plan includes a comprehensive study of deep learning architectures on this dataset. Furthermore, while we have used the Dawid-Skene vote aggregation strategy, it is possible that newer vote aggregation methods can incorporate annotator statistics that may provide a more consistent ground truth, and we leave this for further research. Towards this end, we also provide annotator statistics from CrowdFlower along with the DAiSEE dataset.

References

[1] Affdex. <http://www.affectiva.com/solutions/affdex/>, [Online; accessed 13-March-2016].

[2] T. Baltru, P. Robinson, L.-P. Morency, et al. Openface: an open source facial behavior analysis toolkit. In *2016*

IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1–10. IEEE, 2016.

[3] A. Baylari and G. A. Montazer. Design a personalized e-learning system based on item response theory and artificial neural network approach. *Expert Systems with Applications*, 36(4):8013–8021, 2009.

[4] A. Burke. Crowdsourcing scientific progress: how crowdflower’s hordes help harvard researchers study tb. *Forbes*, October, 16, 2011.

[5] R. A. Calvo and S. D’Mello. Affect detection: An interdisciplinary review of models, methods, and their applications. *Affective Computing, IEEE Transactions on*, 1(1):18–37, 2010.

[6] F. Castro, A. Vellido, À. Nebot, and F. Mugica. Applying data mining techniques to e-learning problems. In *Evolution of teaching and learning paradigms in intelligent environment*, pages 183–221. Springer, 2007.

[7] O. Culture. Moocs interrupted. http://www.openculture.com/2013/04/10_reasons_you_didnt_complete_a_mooc.html, [Online; accessed 11-March-2016].

[8] A. P. Dawid and A. M. Skene. Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied statistics*, pages 20–28, 1979.

[9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.

[10] P. Eckman. Universal and cultural differences in facial expression of emotion. In *Nebraska symposium on motivation*, volume 19, pages 207–284. University of Nebraska Press Lincoln, 1972.







Figure 6. Average classwise accuracy for each affective state for all classifiers, k -NN, SVM and Random Forests and 3 feature extraction methods, namely, 3D-HOG, 3D-LBP and VGG

- [11] Emotient. <http://www.emotient.com/>, [Online; accessed 13-March-2016].
- [12] Emovu. <http://www.emovu.com/e/>, [Online; accessed 13-March-2016].
- [13] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) chal-




Figure 8. Average 2-class accuracy (classwise) for engagement for all classifiers, k -NN, SVM and Random Forests and 3 feature extraction methods, namely, 3D-HOG, 3D-LBP and VGG

E	3	1	2	4
F	2	2	2	1
C	2	1	2	1
B	2	4	3	1

Figure 9. Labels predicted by using 3D-LBP with Random Forests classifier. E=Engagement; F=Frustration; C=Confusion; B=Boredom

lenge. *International journal of computer vision*, 88(2):303–338, 2010.

- [14] H. Gunes, B. Schuller, M. Pantic, and R. Cowie. Emotion representation, analysis and synthesis in continuous space: A survey. In *Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, pages 827–834. IEEE, 2011.
- [15] H. Han, C. Otto, X. Liu, and A. K. Jain. Demographic estimation from face images: Human vs. machine performance. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(6):1148–1161, 2015.
- [16] W. Hawthorne Effect. https://en.wikipedia.org/wiki/hawthorne_effect, 2016.
- [17] L. He, D. Jiang, L. Yang, E. Pei, P. Wu, and H. Sahli. Multimodal affective dimension prediction using deep bidirectional long short-term memory recurrent neural networks. In *Proceedings of the 5th International Workshop on Audio/Visual Emotion Challenge*, pages 73–80. ACM, 2015.
- [18] J. Hernandez, Z. Liu, G. Hulten, D. DeBarr, K. Krum, and Z. Zhang. Measuring the engagement level of tv viewers. In *Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on*, pages 1–7. IEEE, 2013.
- [19] M.-J. Huang, H.-S. Huang, and M.-Y. Chen. Constructing a personalized e-learning system based on genetic algorithm and case-based reasoning approach. *Expert Systems with Applications*, 33(3):551–564, 2007.
- [20] P. G. Ipeirotis, F. Provost, V. S. Sheng, and J. Wang. Repeated labeling using multiple noisy labelers. *Data Mining*

- and *Knowledge Discovery*, 28(2):402–441, 2014.
- [21] K. Jordan. Initial trends in enrolment and completion of massive open online courses. *The International Review of Research in Open and Distributed Learning*, 15(1), 2014.
 - [22] A. Klaser, M. Marszałek, and C. Schmid. A spatio-temporal descriptor based on 3d-gradients. In *BMVC 2008-19th British Machine Vision Conference*, pages 275–1. British Machine Vision Association, 2008.
 - [23] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014*, pages 740–755. Springer, 2014.
 - [24] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*, pages 94–101. IEEE, 2010.
 - [25] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. *Proceedings of the British Machine Vision*, 1(3):6, 2015.
 - [26] M. Pivec, C. Trummer, and J. Pripfl. Eye-tracking adaptable e-learning and content authoring support. *Informatica*, 30(1), 2006.
 - [27] R. Rajendran. *Enriching the Student Model in an Intelligent Tutoring System*. PhD thesis, The IITB-Monash Research Academy, 2014.
 - [28] J. Rehg, G. Abowd, A. Rozga, M. Romero, M. Clements, S. Sclaroff, I. Essa, O. Ousley, Y. Li, C. Kim, et al. Decoding children’s social behavior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3414–3421, 2013.
 - [29] V. S. Sheng, F. Provost, and P. G. Ipeirotis. Get another label? improving data quality and data mining using multiple, noisy labelers. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 614–622. ACM, 2008.
 - [30] SightCorp. <http://www.sightcorp.com/>, [Online; accessed 13-March-2016].
 - [31] J. Tao and T. Tan. Affective computing: A review. In *Affective computing and intelligent interaction*, pages 981–995. Springer, 2005.
 - [32] Y.-I. Tian, T. Kanade, and J. F. Cohn. Recognizing action units for facial expression analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(2):97–115, 2001.
 - [33] M. Valstar, B. Schuller, K. Smith, T. Almaev, F. Eyben, J. Krajewski, R. Cowie, and M. Pantic. Avec 2014: 3d dimensional affect and depression recognition challenge. In *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge*, pages 3–10. ACM, 2014.
 - [34] G. Van Horn, S. Branson, R. Farrell, S. Haber, J. Barry, P. Ipeirotis, P. Perona, and S. Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 595–604, 2015.
 - [35] A. Vedaldi and B. Fulkerson. Vlfeat: An open and portable library of computer vision algorithms. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1469–1472. ACM, 2010.
 - [36] P. Welinder and P. Perona. Online crowdsourcing: rating annotators and obtaining cost-effective labels. 2010.
 - [37] J. Whitehill, Z. Serpell, Y.-C. Lin, A. Foster, and J. R. Movellan. The faces of engagement: Automatic recognition of student engagement from facial expressions. *Affective Computing, IEEE Transactions on*, 5(1):86–98, 2014.
 - [38] P. Yang, Q. Liu, and D. N. Metaxas. Boosting coded dynamic features for facial action units and facial expression recognition. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*, pages 1–6. IEEE, 2007.
 - [39] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang. A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(1):39–58, 2009.
 - [40] G. Zhao and M. Pietikainen. Dynamic texture recognition using local binary patterns with an application to facial expressions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(6):915–928, 2007.
 - [41] D. Zhou, S. Basu, Y. Mao, and J. C. Platt. Learning from the wisdom of crowds by minimax entropy. In *Advances in Neural Information Processing Systems*, pages 2195–2203, 2012.